arXiv:1610.05714v1 [math.CO] 18 Oct 2016

MINIMIZING THE NUMBER OF INDEPENDENT SETS IN TRIANGLE-FREE REGULAR GRAPHS

JONATHAN CUTLER AND A. J. RADCLIFFE

Abstract. Recently, Davies, Jenssen, Perkins, and Roberts gave a very nice proof of the result (due, in various parts, to Kahn, Galvin-Tetali, and Zhao) that the independence polynomial of a d-regular graph is maximized by disjoint copies of $K_{d,d}$. Their proof uses linear programming bounds on the distribution of a cleverly chosen random variable. In this paper, we use this method to give lower bounds on the independence polynomial of regular graphs. We also give new bounds on the number of independent sets in triangle-free regular graphs.

1. Introduction

Extremal problems involving the number of substructures of a graph of a given type have popped up in quite a few different contexts of late. One of the best known such results is due to Kahn [7] and Zhao [10]. We let $\mathrm{Ind}(G)$ be the set of independent sets in G. Their theorem bounds $\mathrm{ind}(G) = |\mathrm{Ind}(G)|$ for regular graphs.

Theorem 1 (Kahn, Zhao)**.** *If G is a d-regular graph on n vertices, then*

$$\mathrm{ind}(G)^{1/n} \leq \mathrm{ind}(K_{d,d})^{1/2d}.$$

One source for questions of this type is the field of statistical mechanics. For instance, the *hard-core model* on a graph G is a probability distribution on the independent sets of G in which a independent set I is chosen with probability proportional to $\lambda^{|I|}$. Here $\lambda > 0$ is a parameter called the *fugacity*. The normalizing factor is

$$P_G(\lambda) = \sum_{I \in \mathrm{Ind}(G)} \lambda^{|I|},$$

known to graph theorists as the *independence polynomial of G* and to statistical physicists as the *partition function* of this hard-core model.

Kahn [7] in fact proved the analogue of Theorem 1 for the independence polynomial of bipartite graphs with fugacity $\lambda \geq 1$, i.e.,

$$P_G(\lambda)^{1/n} \leq P_{K_{d,d}}(\lambda)^{1/2d}.$$

Galvin and Tetali [5] extended Kahn's result to cover the case $0 < \lambda < 1$. Finally, Zhao [10] proved the full theorem using a clever lifting argument.

More recently, Davies, Jenssen, Perkins, and Roberts [3] gave an independent proof introducing an audacious new approach utilizing linear programming. Following Davies et al., we will derive bounds on $P_G(\lambda)$ by considering the *occupancy fraction*, denoted $\alpha_G(\lambda)$. This is the expected fraction of vertices of G belonging to a random independent set chosen according the hard-core model. More explicitly,

$$\alpha_G(\lambda) = \frac{1}{n} \frac{\sum_{I \in \mathrm{Ind}(G)} |I| \lambda^{|I|}}{P_G(\lambda)} = \frac{1}{n} \frac{\lambda P_G'(\lambda)}{P_G(\lambda)}.$$

Davies et al. [3] proved the following.

Date: October 19, 2016.

The first author was sponsored by the National Security Agency under Grant H98230-15-1-0016. The United States Government is authorized to reproduce and distribute reprints notwithstanding any copyright notation herein. The second author was sponsored by the Simons Foundation under Grant 429383.

Theorem 2 (Davies, Jenssen, Perkins, Roberts)**.** *For all d-regular graphs G and all $\lambda > 0$, it is the case that*

$$\alpha_G(\lambda) \leq \alpha_{K_{d,d}}(\lambda).$$

Because $\alpha_G(\lambda)$ is essentially the logarithmic derivative of $P_G(\lambda)$, this is a strengthening of Theorem 1. Corollary 14 below shows how to use the occupancy fraction to bound the independence polynomial.

In this paper, we investigate lower bound analogues of Theorem 2. In Section 2, we prove that for any d-regular graph, the occupancy fraction is bounded below by that of K_{d+1}. This is a stronger form of the corresponding result for $\mathrm{ind}(G)$ proved by the authors in [2]. In later sections, we discuss a problem raised by Kahn [6], that of giving a lower bound on $P_G(\lambda)$ for d-regular triangle-free graphs. We use the same occupancy fraction approach to give bounds in this case.

When λ is large, the hard-core model is biased strongly towards large independent sets. Indeed,

$$\lim_{\lambda\to\infty} \alpha_G(\lambda) = \frac{\alpha(G)}{n}, \tag{1}$$

where $\alpha(G)$ is the *independence number of* G, and the ratio $\alpha(G)/n$ is the *independence ratio of* G. Our methods yield a weak lower bound on the independence ratio of regular triangle-free graphs, which we present in Section 3. In Section 4, we focus on triangle-free cubic graphs. Here we are able to give a bound that is relatively good when $\lambda = 1$. We conjecture that the Petersen graph is extremal, but this cannot be true for all λ since the independence ratio for triangle-free cubic graphs is minimized by $\mathrm{GP}(7, 2)$, a generalized Petersen graph (see Staton [9]).

2. Lower bounds for the hard-core model on regular graphs

The following lower bound on the occupancy fraction is the main theorem of this section. Its proof appears at the end of this section after a number of lemmas concerning the linear programming approach.

Theorem 3. *If G is a d-regular graph on n vertices and $\lambda > 0$, then*

$$\alpha_G(\lambda) \geq \alpha_{K_{d+1}}(\lambda).$$

As a consequence, we have

$$P_G(\lambda)^{1/n} \geq P_{K_{d+1}}(\lambda)^{1/(d+1)}.$$

Equality is, in both cases, only achieved for G a disjoint union K_{d+1}s.

Following Davies et al., we will consider, for each vertex in V, the probability that it belongs to, and the probability that it is covered by, a randomly chosen independent set. To be explicit, if I is an independent set, we say that $v \in V(G)$ is *occupied* if $v \in I$ and *uncovered* if $I \cap N(v) = \emptyset$. If I is distributed according to the hard-core model, we write p_v for $\mathbb{P}(v \in I)$ and q_v for $\mathbb{P}(v \text{ is uncovered})$. Note that both p_v and q_v are functions of λ. Also, it is the case that $p_v \leq q_v$ since $\{I \in \mathrm{Ind}(G) : v \text{ is occupied}\} \subseteq \{I \in \mathrm{Ind}(G) : v \text{ is uncovered}\}$.

Lemma 4. *In the hard-core model on G with fugacity $\lambda > 0$, we have*
1) $p_v = \frac{\lambda}{1+\lambda} q_v$, and
2) $\alpha_G(\lambda) = \frac{1}{n} \sum_{v\in V(G)} p_v$.

Proof. For 1), note that the conditional probability that v is occupied given that it is uncovered is $\lambda/(1+\lambda)$. For the second part, simply write $|I| = \sum_{v\in V(G)} \mathbb{1}(v \in I)$ and take expectations. □

We will prove Theorem 3 by computing the occupancy fraction in two different ways, each based on the neighborhood of a uniformly randomly chosen vertex v in G. In particular, we record the external influence of I on $N(v)$.

Definition. Let $I \in \mathrm{Ind}(G)$ be chosen according to the hard-core model, and v be chosen uniformly from $V(G)$, independently of I. We define random variables

$$U = U(v, I) = N(v) \setminus N(I \setminus N(v)), \qquad \text{and} \qquad H = H(v, I) = G[U].$$

Thus, U is the subset of the neighborhood of v which is not covered by any vertex of I outside of $N(v)$.

We define this triple because, conditioning on $I \setminus N(v)$, we have that $I \cap N(v) \subseteq U$ and, moreover, $I \cap N(v)$ is distributed according to the hard-core model on H with fugacity λ.

Lemma 5 (Davies et al. [3])**.** *In the hard-core model on a d-regular graph G with fugacity $\lambda > 0$ we have, with the notation above,*

1) $\alpha_G(\lambda) = \dfrac{\lambda}{1+\lambda}\mathbb{E}\left(\dfrac{1}{P_H(\lambda)}\right)$*, and also*

2) $\alpha_G(\lambda) = \dfrac{\lambda}{d}\,\mathbb{E}\left(\dfrac{P_H'(\lambda)}{P_H(\lambda)}\right)$.

Proof. Since v is uniformly distributed on $V(G)$, Lemma 4 yields

$$\alpha_G(\lambda) = \mathbb{E}(p_v) = \frac{\lambda}{1+\lambda}\mathbb{E}(q_v) = \frac{\lambda}{1+\lambda}\mathbb{E}\left(\frac{1}{P_H(\lambda)}\right).$$

The final equality follows since v is uncovered precisely if $I \cap N(v) = \emptyset$, and $I \cap N(v)$ is distributed according to the hard-core model on H.

For the second part, we pick a random vertex of $V(G)$ by picking a uniformly random neighbor, say u, of v. Since G is regular, u is also uniformly distributed on $V(G)$. Thus,

$$\alpha_G(\lambda) = \mathbb{E}\,(p_u) = \frac{1}{d}\,\mathbb{E}\left(\frac{\lambda P_H'(\lambda)}{P_H(\lambda)}\right),$$

since $\mathbb{E}(\lambda P_H'(\lambda)/P_H(\lambda))$ is the expected number of occupied neighbors of v. □

We will now find the minimum value of $\mathbb{E}(1/P_H(\lambda))$, where the distribution of H is no longer tied to that arising from some d-regular graph. Instead, the distribution of H will merely have to satisfy the very limited condition that the two expressions for α in Lemma 5 agree. We let $\mathcal{H}$ be the set of all graphs on at most d vertices, including the graph with empty vertex set. We say a random variable H with values in $\mathcal{H}$ is *neighborly* if

$$\frac{\lambda}{1+\lambda}\mathbb{E}\left(\frac{1}{P_H(\lambda)}\right) = \frac{\lambda}{d}\,\mathbb{E}\left(\frac{P_H'(\lambda)}{P_H(\lambda)}\right).$$

Now we define

$$\alpha_* = \frac{\lambda}{1+\lambda}\inf\left\{\mathbb{E}\left(\frac{1}{P_H(\lambda)}\right) : H \text{ is a neighborly probability distribution on } \mathcal{H}\right\}.$$

This minimum is the optimal value of a linear program where the variables are the probabilities that the distribution of H assigns to graphs in $\mathcal{H}$. We write p_H for these probabilities, and also set

$$a_H = \frac{1}{P_H(\lambda)} \qquad \text{and} \qquad b_H = \frac{(1+\lambda)P_H'(\lambda)}{dP_H(\lambda)}.$$

In standard form, the linear program is the following, which we refer to as $\mathrm{LP}(d,\lambda)$.

$$\begin{aligned}
\alpha_* = \min \frac{\lambda}{2(1+\lambda)} \sum_{H\in\mathcal{H}} p_H(a_H + b_H) \quad &\text{subject to} \\
\sum_{H\in\mathcal{H}} p_H &= 1, \\
\sum_{H\in\mathcal{H}} p_H(a_H - b_H) &= 0, \\
p_H \geq 0 \quad &\text{for all } H \in \mathcal{H}.
\end{aligned}$$

A number of times in this paper we use some basic facts about the solutions to linear programs and their duals. Firstly, if we can find feasible solutions to a program and its dual with matching objective values, both solutions are optimal. Secondly, there is a more subtle equivalent criterion for simultaneous optimality called *complementary slackness*. A pair of feasible solutions (one for the primal and one for the dual) satisfy

complementary slackness if, for every matching pair of variable and constraint, either the constraint is tight or the variable is zero (or both). See the text of Chvátal [1] for more details.

We will compute the solution to $\mathrm{LP}(d,\lambda)$ by exhibiting a solution to the primal and a solution to the dual which have matching objective values. We will also verify uniqueness by using complementary slackness. We start by writing down the dual problem. It has two unbounded variables, A and B, corresponding to the equality constraints in the original program and an inequality corresponding to each variable.

$$\begin{aligned}\alpha_* = \max \frac{\lambda}{2(1+\lambda)}A \quad &\text{subject to}\\ A + B(a_H - b_H) \le a_H + b_H \quad &\text{for all } H \in \mathcal{H}.\end{aligned}$$

We will exhibit a dual feasible (A,B) whose dual objective value agrees with the primal objective value for the distribution (p_H) arising from taking $G = K_{d+1}$. To that end, we prove a lemma that describes the (A,B) that satisfy certain of the dual constraints with equality.

Lemma 6. *Suppose that $A = A_K, B = B_K \in \mathbb{R}$ are such that the dual constraints $A+B(a_H-b_H) \ge a_H+b_H$ are satisfied with equality for H being the graph with no vertices (denoted $\emptyset$) and also for $H = K \ne \emptyset$. Then*

$$A_K = \frac{2b_K}{1-a_K+b_K} = \frac{2}{1+\frac{\lambda d}{2(1+\lambda)}\frac{p'(K)}{\mu(K)}} \qquad \text{and} \qquad B_K = 1 - A_K,$$

where $p'(K) = \mathbb{P}(I \ne \emptyset)$ and $\mu(K) = \mathbb{E}|I|$ for I distributed according to $\mathrm{HC}_K(\lambda)$.

Proof. Since $a_\emptyset = 1$ and $b_\emptyset = 0$, one of the equations that A and B satisfy is that $A + B = 1$. Substituting into the constraint corresponding to $H = K$ gives

$$A_K = \frac{2b_K}{1-a_K+b_K} = \frac{2b_K}{p'(K)+b_K} = \frac{2}{1+\frac{p'(K)}{b_K}} = \frac{2}{1+\frac{\lambda d}{2(1+\lambda)}\frac{p'(K)}{\mu(K)}}.$$ □

Lemma 7. *With the notation of the previous lemma, if H and K are graphs on at least one and at most d vertices, then the following are equivalent.*

(1) $A_K + B_K(a_H - b_H) \le a_H + b_H$.
(2) $A_K \le A_H$.
(3) $\dfrac{p'(K)}{\mu(K)} \ge \dfrac{p'(H)}{\mu(H)}$.

Proof. For the first equivalence, note that the following are equivalent.

$$\begin{aligned} A_K + B_K(a_H - b_H) &\le a_H + b_H\\ A_K + (1-A_K)(a_H - b_H) &\le a_H + b_H\\ A_K(1 - a_H + b_H) &\le 2b_H\\ A_K &\le A_H.\end{aligned}$$

The second equivalence comes from the fact that $A_K = 1/\big(1+\frac{\lambda d}{2(1+\lambda)}\frac{p'(K)}{\mu(K)}\big)$ is a strictly decreasing function of $p'(K)/\mu(K)$. □

Now we claim that $A = A_{K_d} = 2(1+\lambda)/(1+(d+1)\lambda)$, $B = B_{K_d} = ((d-1)\lambda-1)/(1+(d+1)\lambda)$ are dual feasible and give the same dual objective value as arises in the primal problem from taking $(p_C) = \big(p_C^{K_{d+1}}\big)$, the probability distribution arising from the graph K_{d+1}. Clearly this choice of A, B gives

$$\alpha = \frac{\lambda}{2(1+\lambda)}A_{K_d} = \frac{\lambda}{1+(d+1)\lambda} = \alpha(K_{d+1}).$$

On the other hand all other dual constraints are satisfied. By Lemma 7, for all H with between 1 and d vertices, $A_K + B_K(a_H - b_H) \le a_H + b_H$, since

$$\frac{p'(K_d)}{\mu(K_d)} = 1 \ge \frac{p'(H)}{\mu(H)}.$$

Thus the minimum value of α is achieved for $G = K_{d+1}$. To prove uniqueness here we need only observe that the last inequality is only tight when H is also complete. Thus, by complementary slackness, no distribution (p_H) can be extremal unless it is supported on complete graphs H (and the zero vertex graph). In particular no graph G can be extremal unless p_G is supported on these. The following lemma characterizes such graphs.

Lemma 8. *If G is d-regular, and for all independent $I \subseteq V(G)$ and all $v \in V$, we have $H = H(v, I)$ complete (or $\emptyset$), then G is a disjoint union of K_{d+1}s.*

Proof. Suppose G is not a disjoint union of K_{d+1}s. Then there exists a vertex v with non-adjacent neighbors u, w. Set $I = \{u, w\}$. Then $H[\{u, w\}] = E_2$ and in particular H is neither $\emptyset$ nor complete. □

Proof of Theorem 3. Given a d-regular graph G, consider the random variables U and H defined just before Lemma 5. Clearly, H is neighborly and thus

$$\alpha_G(\lambda) \geq \alpha_* = \alpha_{K_{d+1}}(\lambda).$$

For the bound on the independence polynomial, note that

$$\begin{aligned}
\log P_G(\lambda) &= \int_0^\lambda \frac{P_G'(t)}{P_G(t)} dt \\
&= n \int_0^\lambda \frac{\alpha_G(t)}{t} dt \\
&\geq n \int_0^\lambda \frac{\alpha_{K_{d+1}}(t)}{t} dt \\
&= \frac{n}{d+1} \log P_{K_{d+1}}(\lambda).
\end{aligned}$$

□

3. Minimizing the hard-core model for triangle-free d-regular graphs

As Davies et al. [3] noted, the linear programming approach is simpler when G is triangle-free since the graph induced on $N(v)$, for any v, is empty. Where before we had to define both U and $H = G[U]$, now U, the set of uncovered neighbors of v, always induces an empty graph. Thus, we need only keep track of how many neighbors of v are uncovered. Our approach will be to add further constraints to the linear program, thereby getting a better approximation to the actual minimum. Pick an independent set I according to the hard-core model and a vertex v uniformly at random, independently of I. Let Y be the number of uncovered neighbors of v, so range$(Y) = \{0, 1, 2, \ldots, d\}$. The distribution of Y is specified by the $d + 1$ values $y_0, y_1, y_2, \ldots, y_d$, where $y_i = \mathbb{P}(Y = i)$.

Restricting ourselves to the constraints analogous to those in LP(d, λ) unfortunately gives a rather weak bound on the hard-core model in triangle-free regular graphs. In this section, we add a natural lower bound on y_0 and, in the next, we add a further constraint in the special case of 3-regular graphs. While neither result is best possible, we are able to give reasonably tight bounds.

The constraint we add in the general d-regular case is the simple one that $y_0 \geq \alpha$. This follows from the fact that any vertex in I must necessarily have all of its neighbors uncovered. Since, by Lemma 5,

$$\alpha_G(\lambda) = \frac{\lambda}{d} \mathbb{E}\left(\frac{P_H'(\lambda)}{P_H(\lambda)}\right) \qquad \text{and} \qquad \frac{P_{E_i}'(\lambda)}{P_{E_i}(\lambda)} = \frac{i(1+\lambda)^{i-1}}{(1+\lambda)^i} = \frac{i}{1+\lambda},$$

the constraint will be written

$$y_0 - \sum_{i=1}^{d} \frac{i\lambda}{d(1+\lambda)} y_i \geq 0.$$

So, the linear program, which we will denote $\mathrm{LP}^+(d,\lambda)$, is

$$\begin{aligned}
\text{Minimize} \quad & \sum_{i=1}^{d} i \cdot y_i \\
\text{subject to} \quad \sum_{i=0}^{d} y_i &= 1, \\
\sum_{i=0}^{d} \big(i - d(1+\lambda)^{-i}\big) y_i &= 0, \\
y_0 - \sum_{i=1}^{d} \frac{i\lambda}{d(1+\lambda)} y_i &\geq 0, \\
y_0, y_1, y_2, \ldots, y_d &\geq 0.
\end{aligned}$$

Note that the solution to $\mathrm{LP}^+(d,\lambda)$ gives (up to a constant factor) a lower bound on $\alpha(G)$ for G a triangle-free d-regular graph.

Theorem 9. *For any λ, the solution to $\mathrm{LP}^+(d,\lambda)$ is supported on three values. In particular, if we define $0 = m_{d-1} < m_{d-2} < \cdots < m_1$ by*

$$m_i = \left(\frac{d}{i+1}\right)^{1/i} - 1,$$

then for $m_i \leq \lambda < m_{i-1}$ (where we let $m_0 = \infty$), the only non-zero variables are y_0, y_i, and y_{i+1}, with optimal values

$$\begin{aligned}
y_0^* &= \frac{\lambda(1+(i+1)\lambda)}{(1+\lambda)^{i+1} + \lambda(d+1+(d+i+1)\lambda)}, \\
y_i^* &= \frac{(1+\lambda)((i+1)(1+\lambda)^i - d)}{(1+\lambda)^{i+1} + \lambda(d+1+(d+i+1)\lambda)}, \\
y_{i+1}^* &= \frac{(1+\lambda)^2(d - i(1+\lambda)^{i-1})}{(1+\lambda)^{i+1} + \lambda(d+1+(d+i+1)\lambda)}.
\end{aligned}$$

Proof. Note that the dual program is as follows.

$$\begin{aligned}
\text{Maximize} \quad & S \\
\text{subject to} \quad S - dM - A &\leq 0, \\
S + \Big(i - \frac{d}{(1+\lambda)^i}\Big) M - \frac{i\lambda}{d(1+\lambda)} A &\leq i, \quad \text{for each } 1 \leq i \leq d, \\
A &\geq 0.
\end{aligned}$$

The solution to the dual program, for $m_i \leq \lambda \leq m_{i-1}$ (or in the case $i = 1$, $m_1 \leq \lambda$), is

$$\begin{aligned}
S^* &= \frac{d(1+\lambda)(1+(i+1)\lambda)}{(1+\lambda)^{i+1} + \lambda(d+1+(d+i+1)\lambda)}, \\
M^* &= \frac{(1+\lambda)^{i+2}}{(1+\lambda)^{i+1} + \lambda(d+1+(d+i+1)\lambda)}, \\
A^* &= \frac{d(1+\lambda)[(1+\lambda)^{i+1} - (1+(i+1)\lambda)]}{(1+\lambda)^{i+1} + \lambda(d+1+(d+i+1)\lambda)}.
\end{aligned}$$

To show that we have solved both linear programs, we simply need to show that both assignments are feasible and that they satisfy complementary slackness. We obtained S^*, M^*, and A^* by insisting that the constraints in the dual corresponding to primal variables y_0, y_i, and y_{i+1} were tight. Therefore, it only remains to verify that the y_i^*s are feasible and that $A^* \geq 0$. The latter is clear since $1 + (i+1)\lambda \leq (1+\lambda)^{i+1}$. We obtained the y_i^* by requiring all the constraints in the primal to be tight and all variables other than y_0,

y_i, and y_{i+1} to be zero. We need to check that these three variables are non-negative, but this follows from the fact that $m_i \le \lambda \le m_{i-1}$. □

Corollary 10. *If G is a d-regular triangle-free graph, then*

$$\alpha(G) \ge y_0^* = \frac{\lambda(1+(i+1)\lambda)}{(1+\lambda)^{i+1} + \lambda(d+1+(d+i+1)\lambda)},$$

where $m_i \le \lambda < m_{i-1}$.

Proof. In our solution, the constraint stating $y_0 \ge \alpha$ is tight. □

As a consequence of this corollary, we do get a weak lower bound on the independence ratio of regular triangle-free graphs. As in (1), we have

$$\frac{\alpha(G)}{n} = \lim_{\lambda\to\infty} \alpha_G(\lambda) \ge \lim_{\lambda\to\infty} \frac{\lambda(1+2\lambda)}{(1+\lambda)^2 + \lambda(d+1+(d+2)\lambda)} = \frac{2}{d+3}.$$

The best known lower bound is due to Shearer [8], who proved the following.

Theorem 11 (Shearer)**.** *If G is a triangle-free d-regular graph, then $\alpha(G)/n(G) \ge f(d)$, where $f(d)$ is given by the recurrence*

$$f(0) = 1, \qquad f(d) = \frac{1+(d^2-d)f(d-1)}{d^2+1}.$$

This bound is stronger than ours, but the bounds are reasonably close for small d.

4. Minimizing the hard-core model for triangle-free cubic graphs

In this section, we focus on the special case of cubic graphs. As in the previous section, we add the constraint $y_0 \ge \alpha = \alpha_G(\lambda)$. We also add a constraint that is a lower bound on p_3. Let T_3 be the first three levels of the infinite 3-regular tree (so that T_3 has ten vertices). Also, recall that $N^2(v) = \{x \in V(G) : d(x,v) = 2\}$ and $N^2[v] = \{x \in V(G) : d(x,v) \le 2\}$.

Lemma 12. *If G is a triangle-free cubic graph and Y is the number of uncovered neighbors of a uniformly chosen vertex with respect to an independent I chosen according to the hard-core model, then*

$$\mathbb{P}(Y=3) \ge \frac{(1+\lambda)^3}{P(T_3)}.$$

Proof. We first note that

$$\mathbb{P}(Y=3) = \mathbb{P}(N^2[v] \cap I \subseteq N(v)).$$

This follows from the fact that all of v's neighbors are uncovered if and only if $v \notin I$ and $N^2(v) \cap I = \emptyset$. Let $A = N^2[v] \cap I$ and $W = V(G) \setminus N^2[v]$. We will bound $\mathbb{P}(A \subseteq N(v) \,|\, I \cap W)$. Since we've conditioned on $I \cap W$, we know that A is distributed as the hard-core model on $G' = G[N^2[v] \setminus N(I \cap W)]$, i.e., the graph on $N^2[v]$ after deleting vertices with neighbors in W. Hence,

$$\mathbb{P}(A \subseteq N(v) \,|\, I \cap W) = \frac{(1+\lambda)^3}{P(G')},$$

since there are precisely eight possible values for A and the generating function for their weights is $(1+\lambda)^3$. It only remains to show that, for all λ, we have $P(G') \le P(T_3)$. Note that there is a size-preserving injection from $\mathrm{Ind}(G')$ to $\mathrm{Ind}(T_3)$ and the result follows. □

To specify the constraint $y_0 \ge \alpha$, note that, using Lemma 4, we have

$$\alpha = \frac{1}{n}\sum_{v\in V(G)} p_v = \frac{1}{n}\cdot\frac{\lambda}{1+\lambda}\sum_{v\in V(G)} q_v = \frac{1}{n}\cdot\frac{\lambda}{1+\lambda}\sum_{v\in V(G)}\sum_{j=0}^{3}\frac{y_j}{(1+\lambda)^j} = \frac{\lambda}{1+\lambda}\cdot\mathbb{E}[(1+\lambda)^{-Y}],$$

where the third equality follows from the fact that $\mathbb{P}(v \text{ is uncovered} \,|\, Y = j) = 1/(1+\lambda)^j$. Also, as observed by Davies et al., we can pick a uniformly random vertex of G by first picking a uniformly vertex v and then picking a uniformly random neighbor of v since G is regular. Thus,

$$\mathbb{E}\, Y = \frac{1}{n} \sum_{v \in V(G)} \sum_{u \in N(v)} q_u = 3 \cdot \frac{1+\lambda}{\lambda} \alpha.$$

So, we have the constraint $\mathbb{E}\, Y = 3\mathbb{E}[(1+\lambda)^{-Y}]$. Of course, $\sum_{i=0}^{3} y_i = 1$.

With our added constraints, the linear program becomes the following. We write Λ for $\frac{(1+\lambda)^3}{P(T_3)}$.

$$\begin{aligned}
\text{Minimize} \qquad & y_1 + 2y_2 + 3y_3 \\
\text{subject to} \qquad & \sum_{i=0}^{3} y_i &&= 1, \\
& \sum_{i=0}^{3} \big(i - 3(1+\lambda)^{-i}\big) y_i &&= 0, \\
& y_0 - \sum_{i=1}^{3} \frac{i\lambda}{3(1+\lambda)} y_i &&\geq 0, \\
& y_3 &&\geq \Lambda, \\
& y_0, y_1, y_2, y_3 &&\geq 0.
\end{aligned}$$

Once again, our strategy will be to exhibit values for the y_is together with values for the dual program which satisfy complementary slackness.

The dual program is

$$\begin{aligned}
\text{Maximize} \qquad & S - \Lambda B \\
\text{subject to} \qquad & S - 3M - A &&\leq 0, \\
& S + \Big(1 - \frac{3}{1+\lambda}\Big) M + \frac{\lambda}{3(1+\lambda)} A &&\leq 1, \\
& S + \left(2 - \frac{3}{(1+\lambda)^2}\right) M + \frac{2\lambda}{3(1+\lambda)} A &&\leq 2, \\
& S + \left(3 - \frac{3}{(1+\lambda)^3}\right) M + \frac{3\lambda}{3(1+\lambda)} A - B &&\leq 3, \\
& A, B &&\geq 0.
\end{aligned}$$

The solution to the primal linear program is as follows.

$$\begin{aligned}
y_0^* &= \frac{\lambda(1+2\lambda)}{1+6\lambda+6\lambda^2} + \frac{\lambda^3(1+\lambda)^2}{(1+6\lambda+6\lambda^2)P(T_3)}, \\
y_1^* &= \frac{-1+\lambda+2\lambda^2}{1+6\lambda+6\lambda^2} + \frac{(1+7\lambda+9\lambda^2+\lambda^3)(1+\lambda)^2}{(1+6\lambda+6\lambda^2)P(T_3)}, \\
y_2^* &= \frac{2(1+\lambda)^2}{1+6\lambda+6\lambda^2} - \frac{(2+14\lambda+21\lambda^2+8\lambda^3)(1+\lambda)^2}{(1+6\lambda+6\lambda^2)P(T_3)}, \\
y_3^* &= \frac{(1+\lambda)^3}{P(T_3)}.
\end{aligned} \tag{2}$$

Each of the y_is is non-negative for all $\lambda \geq 0$. This obvious for y_0 and y_3; for y_1 and y_2 we verify that for each, written as a rational function with denominator $(1+6\lambda+6\lambda^2)P(T_3)$, the numerator have all nonnegative

coefficients. The optimal dual solution is

$$\begin{aligned} S^* &= \frac{3(1+\lambda)(1+2\lambda)}{1+6\lambda+6\lambda^2}, \\ M^* &= \frac{(1+\lambda)^3}{1+6\lambda+6\lambda^2}, \\ A^* &= \frac{3\lambda^2(1+\lambda)}{1+6\lambda+6\lambda^2}, \\ B^* &= \frac{3\lambda^2}{1+6\lambda+6\lambda^2}. \end{aligned}$$

To check complementary slackness, we observe that in fact we have equality in all constraints in both linear programs. Hence, complementary slackness yields that these are each optimal solutions to the corresponding programs. This is summarized in the following theorem.

Theorem 13. *If G is a triangle-free cubic graph on n vertices, then the hard-core model on G with fugacity λ satisfies*

$$\alpha \geq y_0^* = \frac{\lambda(1+2\lambda)}{1+6\lambda+6\lambda^2} + \frac{\lambda^3(1+\lambda)^2}{(1+6\lambda+6\lambda^2)P(T_3)}.$$

Proof. The solution α_* of the minimization problem above is attained for the solution (2). Moreover, since one of our constraints is of the form $y_0 \geq \alpha$, and this constraint is satisfied with equality, we have $\alpha_* = y_0^*$. □

Corollary 14. *If G is a triangle-free cubic graph on n vertices, then for any $\lambda_0 \geq 0$,*

$$\frac{1}{n}\log P_G(\lambda_0) \geq \int_0^{\lambda_0} \frac{y_0^*}{\lambda} d\lambda.$$

In particular,

$$\mathrm{ind}(G)^{1/n} \geq 1.538339. \tag{3}$$

Proof. We have, as in the proof of Theorem 3,

$$\log P_G(\lambda) \geq n\int_0^{\lambda} \frac{\alpha_*(t)}{t} dt = n\int_0^{\lambda} \frac{y_0^*}{t} dt.$$

Numerical integration up to $\lambda = 1$ gives the second inequality. □

Unfortunately, in contrast to the result of Davies et al. and our result from Section 2, we do not determine the extremal graph for the occupancy fraction. It should be noted that, in a recent paper, Davies et al. [4] give a lower bound on the independence number of triangle-free graphs of given maximum degree d that is asymptotically correct as $d \to \infty$. For $d = 3$ and $\lambda = 1$, their bound is

$$P_G(1)^{1/n} \geq \exp\left\{\frac{W(3\log 2)^2 + 2W(3\log 2)}{6}\right\} = 1.516712\ldots,$$

where W is the Lambert W function.

Two graphs provide some support that our bound is not far from being optimal. For the Petersen graph, the left hand side of (3) is at most 1.54199, whereas our bound gave 1.538339. One might even be tempted to think that the Petersen graph is the extremal graph for the occupancy fraction for all λ. However, this cannot be true as a result of Staton [9] yields that a related graph, the *generalized Petersen graph* $\mathrm{GP}(7,2)$ (see Figure 1), has a smaller independence ratio and hence, by (1), has smaller occupancy fraction for large λ.

Theorem 15 (Staton)**.** *If G is a triangle-free cubic graph on n vertices, then*

$$\frac{\alpha(G)}{n} \geq \frac{5}{14} = \frac{\alpha(\mathrm{GP}(7,2))}{n(\mathrm{GP}(7,2))}.$$

We therefore make the following conjecture.

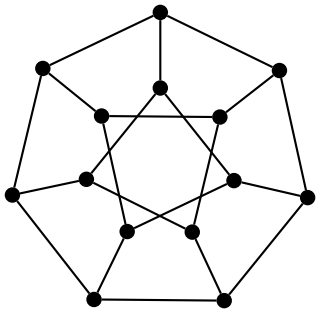

FIGURE 1. The generalized Petersen graph $\mathrm{GP}(7,2)$

Conjecture. If G is a triangle-free cubic graph on n vertices, then

$$P_G(\lambda)^{1/n} \geq \min\left\{P_{\mathrm{GP}(5,2)}(\lambda)^{1/10}, P_{\mathrm{GP}(7,2)}(\lambda)^{1/14}\right\},$$

where $\mathrm{GP}(5,2)$ is the Petersen graph. In particular,

$$\mathrm{ind}(G)^{1/n} \geq \mathrm{ind}(\mathrm{GP}(5,2))^{1/10} = 1.54198\ldots.$$

REFERENCES

1. Vašek Chvátal, *Linear programming*, A Series of Books in the Mathematical Sciences, W. H. Freeman and Company, New York, 1983.
2. Jonathan Cutler and A. J. Radcliffe, *The maximum number of complete subgraphs in a graph with given maximum degree*, J. Combin. Theory Ser. B **104** (2014), 60–71.
3. Ewan Davies, Matthew Jenssen, Will Perkins, and Barnaby Roberts, *Independent sets, matchings, and occupancy fractions*, arXiv: 1508.04675 (2015).
4. ———, *On the average size of independent sets in triangle-free graphs*, arXiv: 1606.01043v1 (2016).
5. David Galvin and Prasad Tetali, *On weighted graph homomorphisms*, Graphs, morphisms and statistical physics, DIMACS Ser. Discrete Math. Theoret. Comput. Sci., vol. 63, Amer. Math. Soc., Providence, RI, 2004, pp. 97–104.
6. Jeff Kahn, personal communication.
7. ———, *An entropy approach to the hard-core model on bipartite graphs*, Combin. Probab. Comput. **10** (2001), no. 3, 219–237.
8. James B. Shearer, *A note on the independence number of triangle-free graphs. II*, J. Combin. Theory Ser. B **53** (1991), no. 2, 300–307.
9. William Staton, *Some Ramsey-type numbers and the independence ratio*, Trans. Amer. Math. Soc. **256** (1979), 353–370.
10. Yufei Zhao, *The number of independent sets in a regular graph*, Combin. Probab. Comput. **19** (2010), no. 2, 315–320.

DEPARTMENT OF MATHEMATICAL SCIENCES, MONTCLAIR STATE UNIVERSITY, MONTCLAIR, NJ
E-mail address: jonathan.cutler@montclair.edu

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF NEBRASKA-LINCOLN, LINCOLN, NE
E-mail address: jamie.radcliffe@unl.edu